

We claim:

1. A method for bypassing a network change by a node in a communication network, the method comprising:
  - 5 pre-determining a recovery path for bypassing a network change that affects communications over a primary path;
  - detecting the network change that affects communications over the primary path; and
  - switching communications from the primary path to the recovery path
  - 10 in order to bypass the network change.
2. The method of claim 1, wherein pre-determining the recovery path for bypassing the network change comprises:
  - establishing as the recovery path a label switched path that bypasses
  - 15 the network change.
3. The method of claim 1, wherein pre-determining the recovery path for bypassing the network change comprises:
  - logically introducing the network change into a routing database; and
  - 20 determining the recovery path based upon a pre-determined path determination scheme.
4. The method of claim 3, wherein the pre-determined path determination scheme comprises a shortest-path-first computation.
- 25 5. The method of claim 1, wherein pre-determining the recovery path for bypassing the network change comprises:
  - installing the recovery path in a forwarding table.
- 30 6. The method of claim 1, wherein detecting the network change that affects communications over the primary path comprises:
  - using a fast liveness protocol to detect the network change.

7. The method of claim 1, wherein the network change comprises a link failure.
- 5 8. The method of claim 1, wherein the network change comprises a node failure.
9. The method of claim 1, wherein the network change comprises a routing change.
- 10 10. The method of claim 1, wherein switching communications from the primary path to the recovery path in order to bypass the network change comprises:  
activating the recovery path.
- 15 11. The method of claim 10, wherein activating the recovery path comprises:  
removing the primary path from a forwarding table.
- 20 12. The method of claim 10, wherein activating the recovery path comprises:  
blocking the primary path in a forwarding table.
13. The method of claim 10, wherein activating the recovery path  
25 comprises:  
marking the recovery path as a higher priority path than the primary path in a forwarding table.
14. The method of claim 1, wherein switching communications from the  
30 primary path to the recovery path in order to bypass the network change comprises:

forwarding all communications from the primary path over the recovery path.

- 15      15.      The method of claim 1, wherein switching communications from the primary path to the recovery path in order to bypass the network change comprises:

forwarding some communications from the primary path over the recovery path based upon a predetermined priority scheme.

- 10      16.      The method of claim 15, wherein the predetermined priority scheme comprises an IP Differentiated Services scheme.

- 15      17.      The method of claim 1, further comprising:  
determining a new primary path.

18.      The method of claim 17, wherein determining the new primary path comprises:  
receiving routing information; and  
computing the new primary path based upon the routing information.

- 20      19.      The method of claim 17, further comprising:  
activating the new primary path.

- 25      20.      The method of claim 19, further comprising:  
switching communications from the recovery path to the new primary path after activating the new primary path.

- 30      21.      The method of claim 19, wherein determining the new primary path and activating the new primary path comprise:  
freezing a forwarding table after switching communications from the primary path to the recovery path;

computing the new primary path while the forwarding table is frozen;  
and

coordinating activation of the new primary path with at least one other  
node in the communication network.

5

22. The method of claim 21, wherein coordinating activation of the new  
primary path with at least one other node in the communication network  
comprises:

using a timer to determine when to activate the new primary path.

10

23. The method of claim 21, wherein coordinating activation of the new  
primary path with at least one other node in the communication network  
comprises:

using a predetermined diffusion mechanism to determine when to  
activate the new primary path.

15

24. The method of claim 21, wherein coordinating activation of the new  
primary path with at least one other node in the communication network  
comprises:

receiving a signal from a master node; and  
activating the new primary path upon receiving the signal from the  
master node.

20

25. The method of claim 21, wherein coordinating activation of the new  
primary path with at least one other node in the communication network  
comprises:

receiving signals from a number of slave nodes;  
determining that the number of slave nodes have completed  
computing new primary paths; and  
activating the new primary path upon determining that the number of  
slave node have completed computing new primary paths.

25

30

26. The method of claim 25, further comprising:  
sending a signal to the number of slave nodes.
27. The method of claim 17, further comprising:  
5 computing a new recovery path to protect the new primary path.
28. The method of claim 19, further comprising:  
computing a new recovery path after activating the new primary path.

29. A device for bypassing a network change in a communication network,  
the device comprising:

recovery path logic operably coupled to pre-determine a recovery path  
for bypassing a network change that affects communications over a primary

5 path;

detection logic operably coupled to detect the network change that  
affects communications over the primary path; and

switching logic operably coupled to switch communications from the  
primary path to the recovery path in order to bypass the network change.

10

30. The device of claim 29, wherein the recovery path logic is operably  
coupled to establish as the recovery path a label switched path that bypasses  
the network change.

15 31. The device of claim 29, wherein the recovery path logic is operably  
coupled to logically introduce the network change into a routing database and  
determine the recovery path based upon a pre-determined path  
determination scheme.

20 32. The device of claim 31, wherein the pre-determined path determination  
scheme comprises a shortest-path-first computation.

33. The device of claim 29, wherein the recovery path logic is operably  
coupled to install the recovery path in a forwarding table.

25

34. The device of claim 29, wherein the detection logic is operably coupled  
to use a fast liveness protocol to detect the network change.

35. The device of claim 29, wherein the network change comprises a link  
30 failure.

36. The device of claim 29, wherein the network change comprises a node failure.

37. The device of claim 29, wherein the network change comprises a  
5 routing change.

38. The device of claim 29, wherein the switching logic is operably coupled to activate the recovery path in order to switch communications from the primary path to the recovery path.  
10

39. The device of claim 38, wherein the switching logic is operably coupled to remove the primary path from a forwarding table in order to activate the recovery path.

40. The device of claim 38, wherein the switching logic is operably coupled to block the primary path in a forwarding table in order to activate the recovery path.  
15

41. The device of claim 38, wherein the switching logic is operably coupled to mark the recovery path as a higher priority path than the primary path in a forwarding table in order to activate the recovery path.  
20

42. The device of claim 29, wherein the switching logic is operably coupled to forward all communications from the primary path over the recovery path.  
25

43. The device of claim 29, wherein the switching logic is operably coupled to forward some communications from the primary path over the recovery path based upon a predetermined priority scheme.

44. The device of claim 43, wherein the predetermined priority scheme comprises an IP Differentiated Services scheme.  
30

45. The device of claim 29, further comprising:  
reconvergence logic operably coupled to determine a new primary  
path.

5 46. The device of claim 45, wherein the reconvergence logic is operably  
coupled to receive routing information and compute the new primary path  
based upon the routing information.

47. The device of claim 45, wherein the reconvergence logic is operably  
10 coupled to activate the new primary path.

48. The device of claim 47, wherein the switching logic is operably coupled  
to switch communications from the recovery path to the new primary path  
upon activation of the new primary path.

15 49. The device of claim 47, wherein the reconvergence logic is operably  
coupled to freeze a forwarding table during computation of the new primary  
path and coordinate activation of the new primary path with at least one  
other node in the communication network.

20 50. The device of claim 49, wherein the reconvergence logic is operably  
coupled to use a timer to determine when to activate the new primary path.

51. The device of claim 49, wherein the reconvergence logic is operably  
25 coupled to use a predetermined diffusion mechanism to determine when to  
activate the new primary path.

52. The device of claim 49, wherein the reconvergence logic is operably  
coupled to receive a signal from a master node and activate the new primary  
30 path upon receiving the signal from the master node.



53. The device of claim 49, wherein the reconvergence logic is operably coupled to activate the new primary path upon determining that a number of slave nodes have completed computing new primary paths based upon signals received from the number of slave nodes.

5

54. The device of claim 53, wherein the reconvergence logic is operably coupled to send a signal to the number of slave nodes upon determining that the number of slave nodes have completed computing new primary paths.

10 55. The device of claim 45, wherein the recovery logic is operably coupled to compute a new recovery path to protect the new primary path.

56. The device of claim 47, wherein the recovery logic is operably coupled to compute a new recovery path after activation of the new primary path.

57. A computer program for programming a computer system to bypass a network change in a communication network, the computer program comprising:

recovery path logic programmed to pre-determine a recovery path for  
5 bypassing a network change that affects communications over a primary path;

detection logic programmed to detect the network change that affects communications over the primary path; and

switching logic programmed to switch communications from the  
10 primary path to the recovery path in order to bypass the network change.

58. The computer program of claim 57, wherein the recovery path logic is programmed to establish as the recovery path a label switched path that bypasses the network change.

15

59. The computer program of claim 57, wherein the recovery path logic is programmed to logically introduce the network change into a routing database and determine the recovery path based upon a pre-determined path determination scheme.

20

60. The computer program of claim 59, wherein the pre-determined path determination scheme comprises a shortest-path-first computation.

61. The computer program of claim 57, wherein the recovery path logic is  
25 programmed to install the recovery path in a forwarding table.

62. The computer program of claim 57, wherein the detection logic is programmed to use a fast liveness protocol to detect the network change.

30 63. The computer program of claim 57, wherein the network change comprises a link failure.

64. The computer program of claim 57, wherein the network change comprises a node failure.

65. The computer program of claim 57, wherein the network change  
5 comprises a routing change.

66. The computer program of claim 57, wherein the switching logic is  
programmed to activate the recovery path in order to switch communications  
from the primary path to the recovery path.

10

67. The computer program of claim 66, wherein the switching logic is  
programmed to remove the primary path from a forwarding table in order to  
activate the recovery path.

68. The computer program of claim 66, wherein the switching logic is  
15 programmed to block the primary path in a forwarding table in order to  
activate the recovery path.

69. The computer program of claim 66, wherein the switching logic is  
20 programmed to mark the recovery path as a higher priority path than the  
primary path in a forwarding table in order to activate the recovery path.

70. The computer program of claim 57, wherein the switching logic is  
programmed to forward all communications from the primary path over the  
25 recovery path.

71. The computer program of claim 57, wherein the switching logic is  
programmed to forward some communications from the primary path over  
the recovery path based upon a predetermined priority scheme.

30

72. The computer program of claim 71, wherein the predetermined  
priority scheme comprises an IP Differentiated Services scheme.

73. The computer program of claim 57, further comprising:  
reconvergence logic programmed to determine a new primary path.

5 74. The computer program of claim 73, wherein the reconvergence logic is  
programmed to receive routing information and compute the new primary  
path based upon the routing information.

75. The computer program of claim 73, wherein the reconvergence logic is  
10 programmed to activate the new primary path.

76. The computer program of claim 75, wherein the switching logic is  
programmed to switch communications from the recovery path to the new  
primary path upon activation of the new primary path.

15 77. The computer program of claim 75, wherein the reconvergence logic is  
programmed to freeze a forwarding table during computation of the new  
primary path and coordinate activation of the new primary path with at least  
one other node in the communication network.

20 78. The computer program of claim 77, wherein the reconvergence logic is  
programmed to use a timer to determine when to activate the new primary  
path.

25 79. The computer program of claim 77, wherein the reconvergence logic is  
programmed to use a predetermined diffusion mechanism to determine when  
to activate the new primary path.

80. The computer program of claim 77, wherein the reconvergence logic is  
30 programmed to receive a signal from a master node and activate the new  
primary path upon receiving the signal from the master node.

81. The computer program of claim 77, wherein the reconvergence logic is programmed to activate the new primary path upon determining that a number of slave nodes have completed computing new primary paths based upon signals received from the number of slave nodes.

5

82. The computer program of claim 81, wherein the reconvergence logic is programmed to send a signal to the number of slave nodes upon determining that the number of slave nodes have completed computing new primary paths.

10

83. The computer program of claim 73, wherein the recovery logic is programmed to compute a new recovery path to protect the new primary path.

15

84. The computer program of claim 75, wherein the recovery logic is programmed to compute a new recovery path after activation of the new primary path.

20

85. The computer program of claim 57 embodied in a computer readable medium.

86. The computer program of claim 57 embodied in a data signal.

87. A communication system comprising a plurality of interconnected communication nodes, wherein primary paths are established for forwarding information and recovery paths are pre-computed for bypassing potential primary path failures.

5

88. The communication system of claim 87, wherein communications are switched from a primary path to a recovery path in order to bypass a network change.

10 89. The communication system of claim 88, wherein new primary paths are determined after communications are switched from the primary path to the recovery path, and communications are switched from the recovery path to a new primary path.

15 90. The communication system of claim 89, wherein each communication node freezes a forwarding table before determining new primary paths.

20 91. The communication system of claim 89, wherein new recovery paths for protecting the new primary paths are computed before switching communications from the recovery path to the new primary path.

92. The communication system of claim 89, wherein new recovery paths for protecting the new primary paths are computer after switching communications from the recovery path to the new primary path.

93. A method for reconverging routes in a communication network, the method comprising:  
determining that a route change is needed;  
freezing forwarding tables so that a predetermined set of routes is used  
5 during reconvergence; and  
reconverging on a new set of routes while the forwarding tables are frozen.

94. The method of claim 93, further comprising:  
10 activating the new set of routes in a coordinated manner.

95. The method of claim 94, wherein activating the new set of routes in a coordinated manner comprises:  
starting a timer by each of a number of nodes in the communication  
15 network upon determining that reconvergence is needed; and  
activating the new set of routes by each of the number of nodes upon expiration of the timer.

96. The method of claim 94, wherein activating the new set of routes in a  
20 coordinated manner comprises:  
using a predetermined diffusion mechanism by each of the number of nodes to determine when reconvergence is complete; and  
activating the new set of routes by each of the number of nodes upon determining that reconvergence is complete.

97. The method of claim 94, wherein activating the new set of routes in a coordinated manner comprises:  
designating one of the number of nodes as a master node and  
designating the remaining nodes as slave nodes;  
30 sending a first signal by each of the slave nodes to the master node upon reconverging on the new set of routes; and





98. A use of a bypass mechanism for bypassing a network change in a communication network, the use comprising:

- 5 using the bypass mechanism to pre-compute a recovery path for bypassing a network change affecting communication over a primary path,  
detect the network change affecting communication over the primary path,  
and switch communications from the primary path to the pre-computed recovery path upon detecting said network change.

99. The use of claim 98, further comprising:

- 10 using the bypass mechanism to compute a new primary path after switching communications from the primary path to the pre-computed recovery path; and

15 using the bypass mechanism to switch communications from the pre-computed recovery path to the new primary path.

100. The use of claim 99, further comprising:

using the bypass mechanism to compute a new recovery path for bypassing a network change affecting communication over the new primary path.